

**Full Mesh LSP and Full Mesh T-LDP Provisioning Between Provider Edge
Routers in Support of Layer-2 and Layer-3 Virtual Private Network services**

Field of the invention

[01] The invention relates to network commissioning in a network management context. In particular to methods of provisioning full mesh Label Switched Path (LSP) connectivity between Provider Edge (PE) routers in support of Layer-2 Virtual Private Network (L2VPN) services such as Virtual Private Wire Services (VPWS) and Virtual Private LAN Service (VPLS).

Background of the invention

[02] In the field of communications, service provider networks are undergoing consolidation due to market pressure to reduce telecommunication service provisioning overheads associated with maintaining multiple networks. In particular, there is a need to address cost/complexity issues related to provisioning Frame Relay (FR) services. Also, there is a need to address high bandwidth overheads incurred during content transport over Asynchronous Transfer Mode (ATM) infrastructure.

[03] Internet Protocol (IP) / MultiProtocol Label Switching (MPLS) -based network provisioning enjoys an extensive deployment largely due to the availability of economic, high speed equipment suited for service-provider-side infrastructure deployment. IP/MPLS networks can support a variety of Layer-2 (L2) technologies; including Ethernet, ATM and FR.

[04] Large organizations (customers) have multitudes of geographically displaced sites, which typically utilize Ethernet Local Area Network (LAN) infrastructure at each site. Traditionally, to provide L2VPN connectivity between these geographically displaced sites, FR services were used for inter-site connectivity between site-specific Customer Edge (CE) nodes because, FR technologies provide traffic differentiation and deterministic content transport. With the advent of IP/MPLS technologies, service

providers can provide similar or better levels of Quality-of-Service (QoS) to the customer.

[05] Service providers seek solutions for connecting their IP/MPLS networks with the customers' Ethernet infrastructure to provide low cost, efficient service offerings to 5 customers, at a throughput comparable to traditional FR Virtual Leased Line (VLL) services and ATM services. However, unlike ATM and FR, plain vanilla Ethernet/IP content transport is performed in accordance with a broadcast/best-effort discipline in accordance with which packets propagate between source and destination nodes non-deterministically, and without guarantees, without necessitating infrastructure 10 redundancy in conveying packets around failed infrastructure.

[06] L2VPN services such as Virtual Private Wire Service (VPWS) and Virtual Private LAN Service (VPLS) can be provisioned over IP/MPLS networks.

[07] VPWS coined by the Internet Engineering Task Force (IETF), generically refers to a L2VPN service that provides an Open Systems Interconnect (OSI) Layer-2 (link 15 layer) point-to-point service (link emulation) connecting two CE nodes, associated with two corresponding customer sites, across a service provider's communications network. VPWS is also known as an Ethernet Virtual Private Line (EVPL) service if the source and destination network nodes are Ethernet based.

[08] To provide content transport in an IP/MPLS network 10, unidirectional Label 20 Switched Paths (LSPs) 12 are created between multiple Label Switching Routers (LSRs) 14 as shown in FIG. 1. In order to provide deterministic content transport and quality-of-service support, source routed LSPs 12 must be established in accordance with a traffic engineering MPLS protocol such as Resource Reservation Protocol-Traffic Engineering (RSVP-TE).

[09] The introduction of MPLS technology not only provides support for 25 deterministic traffic content transport but also provides a migration path for service providers in support of convergence leveraging currently deployed infrastructure. A multitude of transport layer/link layer/physical layer protocols and infrastructure are compatible with MPLS in conveying content. As such, MPLS may be provisioned over 30 IP links 16 themselves provisioned over: an IP/Ethernet infrastructure (GigE), a Packet

Over Synchronous Optical NETwork (SONET) (POS) infrastructure, an Asynchronous Transfer Mode (ATM) infrastructure, etc. No physical layer details of the infrastructure of the communications network 10 are shown in FIG. 1 nor relevant as an abstraction thereof is made in employing MPLS technologies.

- 5 [10] In accordance with customer content traffic differentiation techniques, CEs 22 associated with customer sites 20 connect to corresponding PEs 24 via respective Attachment Circuits (ACs) 26. Each AC 26 may be either a physical or a logical circuit provisioned in accordance with the installed communications network edge infrastructure in the aggregation/distribution portion of the communications network 10.
- 10 10 PEs 24 multiplex customer traffic (30) onto corresponding L2VPN connection(s) 40 (a pseudo wire L2VPN connection is show).

- [11] A Pseudo Wire (PW) 40 refers to an emulated bi-directional point-to-point connection over a packet-switched communications network 10 providing connectivity between two remote network nodes 22 employing any OSI layer-2 technology allowing content traffic arriving an interface 28 of a PE switch 24 to be directed across the core of communications network 10 to a corresponding interface 28 on the corresponding peer PE 24. PW technology gets its importance from what it brings to customers by extending customer services across long distances. PW mechanisms provide emulation of the essential attributes of a selected service through a core communications network
- 15 10 of a different infrastructure (and transport technology). From a customer's point of view, a PW 40 acts as an unshared link or circuit of a particular service – each content frame conveyed by one CE 22 over the PW 40 is received by, and only by, the remote peer CE 22. Content frame forwarding via a PW 40 is not affected by the content frames themselves, rather defined by the end-to-end PW virtual circuit 40 to which the
- 20 25 content frame is submitted for transmission.

- [12] In order to support PW 40 connectivity in an MPLS domain, a targeted LDP session must be established between corresponding PEs 24 enabling the exchange of MPLS labels. Each targeted LDP session includes an LDP signaling link configured with knowledge of remote-peer loopback address information at each end. Only a single targeted LDP session is allowed between any pair of PEs 24. Further information regarding PW and VPLS support is provided in draft-ietf-pwe3-requirements-08.txt,

published by the Internet Engineering Task Force (IETF), specification which is incorporated herein by reference.

[13] Therefore, to support any arbitrary PW 40 in a service provider's communications network 10, a full mesh of bi-directional LSP tunnels 12 and targeted 5 LDP sessions need to be established between participating PE nodes 24. The targeted LDP sessions enable setting up PWs 40 over the full mesh of LSP tunnels 12 which carry the PW content traffic.

[14] Currently, PW-using-LSP connectivity is provisioned manually. The problem is that tens, or even hundreds, of PE nodes 24 are typically employed in a particular 10 service provider network 10 to provide customer services and therefore the number of fully meshed bi-directional LSP tunnels and targeted LDP sessions between N PE 24 nodes is in the order of $N(N-1)$. Therefore manual setup provisioning is lengthy and error-prone. Also, once the full mesh is provisioned, manually adding or removing a PE node 24 is also time consuming and error-prone.

15 [15] A prior art United States Application publication number 2003/0177221 A1 entitled "Resource Allocation Using an Auto-Discovery Mechanism for Provider-Provisioned Layer-2 and Layer-3 Virtual Private Networks" which was published on September 18th, 2003, describes a method in accordance with which, rather than manually configuring VPN tunnels at each PE router, the VPN Capability Discovery 20 Information (VCDI) is "piggy-backed" onto auto-discovery information as an extension to a conventional information distribution protocol, such as Border Gateway Protocol (BGP), Domain Name Service (DNS), and RADIUS. In order to implement the proposed method, auto-discovery protocols have to be extended to include the transmission of the VCDI information. After receiving such information, a tunnel is 25 established by, and between, the appropriate PE nodes based on the information. While the solution has merit, it requires that each PE implement the modified auto-discovery protocol. The requirement suffers from complex implications related to the fact that multi-vendor heterogeneous PE equipment is typically employed in a typical communication network infrastructure and therefore the proposed solution need be 30 adopted by multiple vendors. While the multiple vendors may adopt such a solution eventually, there is no telling how fast the solution may be deployed as all PE

equipment will need upgrading. Further, vendors may see no incentive in upgrading older PE equipment rendering it obsolete and thus preventing service providers from leveraging their existing infrastructure investment.

[16] Therefore there is a need to solve the above mentioned issues in provisioning
5 VPN services over a managed communications network.

Summary of the invention

[17] In accordance with an aspect of the invention, a method of provisioning support
for Virtual Private Network (VPN) services, centrally, in a network management
context is provided. The method includes establishing a full transport mesh of bi-
10 directional Label Switched Paths (LSPs) between a multitude of managed Provider
Edge (PE) communications network nodes in a managed communications network, and
establishing a full signaling mesh of targeted Label Distribution Protocol (LDP)
signaling sessions between the multitude of PE communications network nodes.

[18] In accordance with another aspect of the invention, in establishing the full
15 transport mesh, the method further includes further steps. A group of PE
communications network nodes participating in the managed communications network
is selected. Each selected PE communications network node has a Label Switched
Router (LSR). A multitude of PE communications network node pairs is derived from
the selected group of PE communications network nodes. And, transport LSP
20 commissioning commands are issued to each PE communications network node in the
group to establish an LSP in respect of each corresponding PE communications network
node pair.

[19] In accordance with a further aspect of the invention, prior to issuing LSP
commissioning commands, the method includes further steps. A determination is made
25 as to whether a managed LSPs already exist between pairs of PE communications
network nodes in the group. A further determination is made regarding whether
provisioning parameters associated with the discovered LSPs comply with
corresponding transport mesh provisioning parameters. And, the discovered LSPs are
conditionally included in the transport mesh if compliance exists.

[20] In accordance with a further aspect of the invention, the method further includes steps of commissioning an additional LSP. A pair of PE communications network nodes is selected, and LSP commissioning commands are issued to each PE communications network node in the pair to establish an additional LSP therebetween.

5 [21] In accordance with a further aspect of the invention, the method further includes steps of decommissioning a selected LSP. The two PE communications network nodes corresponding to the ends of the selected LSP are identified. A determination is made regarding whether content is being conveyed via the selected LSP in respect of an actively provisioned VPN service. And, LSP decommissioning commands are

10 selectively issued to the two PE communications network end nodes if no content is being conveyed via the selected LSP.

[22] In accordance with a further aspect of the invention, the method further includes steps adding a PE communications network node to the group. A multitude of PE communications network node pairs is derived, each pair including the additional PE

15 communications network node and one of the PE communications network nodes in the group. LSP commissioning commands are issued to the additional PE communications network node and each PE communications network node in the group, to establish an LSP corresponding to each PE communications network node pair determined.

[23] In accordance with a further aspect of the invention, the method further includes

20 steps removing a PE communications network node from the group. Provisioned LSPs terminating at the PE communications network node to be removed are identified. PE communications network nodes corresponding to the ends of the identified LSPs are also identified. A determination is made regarding whether content is being conveyed via the identified LSPs in respect of at least one actively provisioned VPN service. LSP

25 decommissioning commands are selectively issued to corresponding PE communications network end nodes for each identified LSP if no content is being conveyed therethrough. And, the PE communications network node is excluded from the group if all identified LSPs have been decommissioned.

[24] In accordance with a further aspect of the invention, the method includes further

30 steps in establishing the full signaling mesh. A group of PE communications network

nodes participating in the managed communications network is selected. A multitude of PE communications network node pairs is derived in respect of the selected PE communications network nodes in the group. And, targeted LDP session commissioning commands are issued to each PE communications network node in the
5 group to establish a targeted LDP session in respect of each corresponding PE communications network node pair.

[25] In accordance with a further aspect of the invention, prior to issuing targeted LDP commissioning commands, the method includes further steps. A determination is made as to whether managed targeted LDP sessions already exist between a pairs of PE
10 communications network nodes in the group. And, discovered targeted LDP sessions are included in the signaling mesh.

[26] In accordance with a further aspect of the invention, the method includes further steps in commissioning an additional targeted LDP session. A pair of PE communications network nodes is selected. A corresponding pair of interfaces may
15 further be selected, each interface being associated with a respective PE communication network node in the pair. And, targeted LDP session commissioning commands are issued to each PE communications network node in the pair to establish the additional targeted LDP session therebetween.

[27] In accordance with a further aspect of the invention, the method includes further
20 steps in decommissioning a selected targeted LDP session. The two PE communications network nodes corresponding to the ends of the selected targeted LDP session are determined. A further determination is made as to whether signaling information is being conveyed via the selected targeted LDP session in respect of at least one actively provisioned VPN service. And, targeted LDP session decommissioning commands are
25 selectively issued to the two PE communications network end nodes if no signaling information is being conveyed via the selected targeted LDP session.

[28] In accordance with a further aspect of the invention, the method includes further steps adding a PE communications network node to the group. A multitude of PE communications network node pairs are determined, each pair including the additional
30 PE communications network node and one of the PE communications network nodes in

the group. Targeted LDP session commissioning commands are issued to the additional PE communications network node and each PE communications network node in the group, to establish a targeted LDP session corresponding to each PE communications network node pair determined.

5 [29] In accordance with a further aspect of the invention, the method includes further steps in removing a PE communications network node from the group. Provisioned targeted LDP session terminating at the PE communications network node to be removed are identified. PE communications network nodes corresponding to the ends of the identified targeted LDP sessions are selected. A determination is made whether

10 signaling information is being conveyed via identified targeted LDP sessions in respect of at least one actively provisioned VPN service. Targeted LDP session decommissioning commands are issued to the selected PE communications network end nodes of no signaling information is being conveyed via the identified targeted LDP sessions. And, the PE communications network node is excluded from the group if all

15 identified targeted LDP sessions have been decommissioned.

[30] In accordance with a further aspect of the invention, the method further includes a step of ascribing an identifier to one of: a transport mesh of LSPs, and a signaling mesh of targeted LDP sessions.

20 [31] In accordance with a further aspect of the invention, the method further includes a step of: tracking one of: a transport mesh and a signaling mesh in a network management system repository.

25 [32] In accordance with a further aspect of the invention, a network management system centrally provisioning full mesh MultiProtocol Label Switching (MPLS) connectivity in a managed communications network in support of Virtual Private Network (VPN) service provisioning is provided. The network management system includes a network management system repository tracking managed Provider Edge (PE) communications network nodes in the managed communications network, a full content transport Label Switched Path (LSP) mesh configuration and provisioning means for managing a multitude of content transport LSPs between a selected group of

30 managed PE communications network nodes; and a full targeted Label Distribution

Protocol (LDP) signaling session mesh configuration and provisioning means for managing a multitude of targeted LDP signaling sessions between the selected group of PE communications network nodes.

[33] In accordance with a further aspect of the invention, the full content transport
5 LSP mesh configuration and provisioning means includes a human-machine-interface.

[34] In accordance with a further aspect of the invention, the full targeted LDP signaling session mesh configuration and provisioning means includes a human-machine interface.

[35] In accordance with a further aspect of the invention, the network management
10 system further includes one of: managed PE communications network node filtering means, and managed PE communications network node pair selection means.

[36] In accordance with yet another aspect of the invention, the network management system further includes one of: LSP discovery means, targeted LDP session discovery means, and LSP parameter comparison means.

15 [37] Advantages are derived from the full mesh LSP content transport provisioning and targeted LDP session provisioning effected in a centralized network management context enabling service providers to: assure network resiliency, assure service quality, and provide accounting, in respect of Layer 2 and Layer 3 VPN including, but not limited to, PW and VPLS services irrespective of client side deployed infrastructure
20 while leveraging installed infrastructure in the service provider communications network core. Additional advantages are derived from increased operations management personnel productivity while reducing the probability of human error compared to manual provisioning thereof enabling wide availability of L2 and L3 VPN services.

25 **Brief description of the drawings**

[38] The features and advantages of the invention will become more apparent from the following detailed description of the exemplary embodiment(s) with reference to the attached diagrams wherein:

FIG. 1 is a schematic diagram showing an exemplary bi-directional pseudo-wire provisioned manually over a pair of label switched paths across a communications network;

5 FIG. 2 is a schematic diagram showing, in accordance with an exemplary embodiment of the invention, a plurality of exemplary content transport full meshes of bi-directional label switched paths and an exemplary full mesh of targeted label distribution protocol signaling sessions;

10 FIG. 3 is a schematic diagram of showing, in accordance with an exemplary implementation of the exemplary embodiment of the invention, a content transport label switched path mesh configuration module;

FIG. 4 is a schematic diagram showing, in accordance with the exemplary implementation of the exemplary embodiment of the invention, a state transition diagram corresponding to a content transport label switched path commissioning process;

15 FIG. 5 is a schematic diagram showing, in accordance with another exemplary implementation of the exemplary embodiment of the invention, exemplary details of a content transport label switched path commissioning process;

20 FIG. 6 is a schematic diagram of showing, in accordance with and exemplary implementation of the exemplary embodiment of the invention, a signaling label distribution protocol session mesh configuration module; and

FIG. 7 is a schematic diagram showing, in accordance with the exemplary implementation of the exemplary embodiment of the invention, a state transition diagram corresponding to a signaling label distribution protocol session commissioning process;

25 FIG. 8 is a schematic diagram showing, in accordance with another exemplary implementation of the exemplary embodiment of the invention, exemplary details of a signaling label distribution protocol session commissioning process;

FIG. 9 is a schematic diagram showing Ethernet-based virtual private LAN service interworking over MPLS; and

FIG. 10 is a schematic diagram showing hybrid ATM/Ethernet virtual private LAN service interworking over MPLS.

5 [39] It will be noted that in the attached diagrams like features bear similar labels.

Detailed description of the embodiments

[40] In accordance with an exemplary embodiment of the invention, full mesh Label Switched Path (LSP) provisioning and full mesh targeted Label Distribution Protocol (LDP) session provisioning in a centralized network management context is provided.

10 [41] In a network management context shown in FIG. 2, at least one Network Management System (NMS) 100 is employed to manage network entities, such as, but not limited to: communications network nodes 102 in a managed communications network 104. Communications network management is important particularly in view of complexities associated with typical service provider communications networks,
15 regarding which the following relevant complexities are only exemplary:

[42] Multi-vendor equipment is employed in a typical communications network 104. Multi-vendor equipment also includes multiple equipment types as well multiple versions of the vendor equipment running multiple versions of software/firmware.

20 [43] Communications network nodes 102 may be typically classified by function and the manner in which the communications network nodes 102 are employed as part of the (installed) infrastructure of the communications network 104. It is important to keep in mind that each communications network node 102 may perform a variety of communications network functions. Typically communications network nodes 102 employed in a typical managed service provider network 104, are configurable multi-shelf devices which implement physical transport technologies and content transport protocols on interface cards 106, although monolithic network node devices 102 may
25 also be employed.

[44] Communications network nodes 102 connected to Customer Edge (CE) equipment 108 via physical access links 110 are referred to as Provider Edge (PE) communications network nodes. Communications network nodes 102 not connected directly to CE equipment 108 are known as core Provider (P) communications network 5 nodes 102. It is understood that the flat communications network 104 representation shown in FIG. 2 is exemplary only, in practice communications network nodes 102, may act simultaneously as PE nodes 102 for some services and as core P nodes 102 for other services without any interference existing between the supported services.

[45] It is important to keep in mind that typical communications network nodes 102 10 employed in service provider networks 104 simultaneously support a variety of physical transport technologies and content transport protocols. Some such communications network nodes 102, may also support MPLS content transport by implementing the MPLS protocol over the physical transport equipment installed therein (interface cards 106, switching back-planes, etc.) while other communications network nodes 102, are 15 not aware of the MPLS content transport and are configured to simply switch content conveyed therethrough in accordance with the underlying content transport protocols supported which enable MPLS content transport on neighboring network node 102.

[46] Each Label Switching Router (LSR) 112 is MPLS enabled and may be implemented either as a hardware device or implemented in software as a virtual LSR 20 router instance. Either implementation of the LSR 112 represents a managed entity in the communications network 104.

[47] The above mentioned complexities are addressed via communications network management. The above mentioned NMS 100 has an associated NMS repository 120 tracking managed entity information for the managed infrastructure of the service 25 provider's communications network 104. As shown, the NMS repository 120 may include for example an NMS database without limiting the invention thereto.

[48] In accordance with the exemplary embodiment of the invention, the functionality provided by the NMS 100 is augmented with centralized Layer-2 and Layer-3 VPN provisioning support.

[49] Aside from cable and link setup, depending on the particular implementation, communications network nodes 102 managed in the network management context may have to be assigned Provider Edge (PE) / Provider (P) communications network node roles. Further, quality-of-service policies may have to be uploaded and activated on 5 participating PE and P communications network nodes 102.

[50] In accordance with the exemplary embodiment of the invention, the functionality of the NMS 100 is augmented with centralized full mesh LSP tunnel establishment means. In accordance with an exemplary full mesh LSP tunnel establishment method, operations management personnel selects a group of PE network nodes 102 to provision 10 PW services therebetween for customer sites connected thereto, LSP tunnel commissioning commands are issued from the NMS 100 to each PE 102 in the group to effect the LSP tunnel establishment therebetween.

[51] In selecting the group 312 of PE network nodes 102, operations management personnel, interacting with a Transport LSP Mesh Configuration Module 300 having an 15 exemplary interface schematically shown in FIG. 3, may be provided with a stylized map of the managed communications network 104 by interacting with a Show Network Map button 302. The network map shows iconical representations managed network entities, from which PE network nodes 102 may be selected. In providing operations management personnel with the stylized map, the NMS 100 queries the NMS repository 20 120 for communications network nodes 102 (and logical positions thereof in the network).

[52] The NMS 100 may further query the NMS repository 120 for a list 304 of communications network nodes 102, by interacting with a Get Node List button 306, to aid operations management personnel in selecting PE communications network nodes 25 102.

[53] Alternatively, in selecting the group 312 of PE communications network nodes 102, having retrieved a list of communications network nodes 102, operations management personnel may be provided with filtering means 308 to identify PE communications network nodes 102 for inclusion 310 into the group 312 of PE network 30 nodes 102.

[54] In order for the transport LSP tunnels to be established between the PEs in the group, each PE communications network node 102 must also have LSR functionality. Having LSR functionality may be one of the selection criteria used in filtering as mentioned above (see “Is LSR” filter 308).

5 [55] In ensuring quality-of-service in support of Service Level Agreements (SLA) in accordance with which VPN services are to be offered, LSP content transport tunnels may be commissioned to convey content therein in accordance with quality-of-service parameters 314 including, but not limited to: Committed Information Rate (CIR), Peak Information Rate (PIR), burst size, signaling protocol, tunnel bandwidth, etc. Quality of service parameters 314 may be specified and accepted by interacting with an Update Mesh Parameters button 316.

10 [56] The details of centralized transport LSP tunnel commissioning are described elsewhere, suffice it to mention that, by interacting with a Commission Transport Mesh button 318, each LSR on a corresponding pair of PE network nodes 102 is instructed by the NMS 100 to establish LSP tunnels therebetween relieving operations management personnel of doing so manually. $N^*(N-1)$ unidirectional LSP tunnels are therefore commissioned in respect of each full content transport mesh in providing bi-directional content transport between N PE communications network nodes 102.

15 [57] It is possible for a content transport LSP tunnel to already be commissioned between a particular pair of PE network nodes 102 prior to full mesh LSP tunnel commissioning, the LSP tunnel already conveying content in accordance with the required transport LSP parameters 314. In accordance with an exemplary implementation of the exemplary embodiment of the invention, a process 500 of commissioning 318 an LSP transport mesh shown in FIG. 5, may include determining 20 whether LSP transport tunnels, themselves managed entities in the network management context, already exist between PE communications network node pairs 502 and, subject 508 to compliant LSP transport tunnel parameters, including 510 the existing LSP transport tunnels in the transport mesh being commissioned. An exemplary iterative (504, 512) process 500 is shown. Accordingly, the NMS full 25 transport mesh commissioning means 300, further includes functionality for discovering 506 actively commissioned LSP content transport tunnels, and for determining 508

30

discovered LSP content transport tunnel adherence to a given set of LSP tunnel operational parameters.

[58] In the following, LSP tunnel establishment between PE network node pairs is assumed, although in practice, depending on the signaling protocol employed, the path 5 taken by content conveyed in respect of an LSP tunnel may vary in accordance with communications network conditions experienced thus providing resilience therein.

[59] In accordance with the exemplary embodiment of the invention, the transport mesh of established LSP tunnels between a group 312 of PE communications network nodes represents a managed entity in the network management context. A transport 10 mesh identifier 322 is associated with the transport mesh managed entity for further selection and reference, and a descriptive string 324 may be further associated with the managed transport mesh entity. A status 326 may be ascribed to the managed transport mesh entity and updated based on a multitude of related events. A Get Transport Mesh List button 328 is provided to query the NMS database 120 for transport mesh managed 15 entities.

[60] The Transport LSP Mesh Configuration and Provisioning module 300 may be employed to create different types of LSPs using LDP or RSVP for signaling providing a choice in creating E-LSPs conveying content traffic for multiple Classes-of-Service (CoS) or creating CoS specific L-LSPs.

20 [61] Alternatively, depending on the implementation and/or details of a particular deployment in support of specific VPN services, multiple transport LSP tunnel full meshes may be commissioned between the same group 312 of PE network nodes 102. Multiple transport LSP tunnel full meshes may exemplary be commissioned to support different levels of service in accordance with a corresponding multitude of SLA 25 agreements. Typically service providers provide three levels of service such as “gold”, “silver”, and “bronze”, each level of service specifying different sets of full mesh transport configuration parameters 314 (burst size, bandwidth, etc.)

[62] Individual PE communications network nodes 102 may be added 310 to or removed 311 from a PE group 312, as well additional LSP tunnels may be

commissioned 330 between the PE network nodes 102 in the group 312, or decommissioned 331.

[63] With each PE network node 102 addition 310 to an existing group of N PE network nodes 312, $2*N$ unidirectional LSP tunnels are commissioned between N pairs 5 of PE network nodes: between the newly added PE network node 102 and the existing N PE network nodes 102 in the group 312. With each PE network node 102 removal 311 from an existing group of N PE network nodes 312, $2*(N-1)$ unidirectional LSP tunnels are decommissioned between N-1 pairs of PE network nodes: between the PE network node to be removed and the remaining N-1 PE network nodes 102 in the group 10 312.

[64] The removal 311 of a PE network node 102 from the PE group may be conditional on the non-existence of any VPN services actively commissioned therethrough. The restriction condition may be further limited to VPN services terminating on the PE network node 102 to be removed 311, as the removed PE 102 15 may still convey content in respect of managed VPN services as an intermediary MPLS communications network node 102.

[65] LSP transport tunnels also represent managed entities in the NMS management context and are tracked in the NMS repository 120. Each LSP transport tunnel has an associated identifier 332, a description 334, and an status 336. The NMS repository 120 20 may be queried, dependent on the active network management context, by interacting with a Get Transport LSP List button 338. If the get transport LSP list button 338 is employed while a transport mesh is selected, then only the LSP transport tunnels of the selected transport mesh are displayed, otherwise all LSP transport tunnels are displayed. LSP tunnels commissioned in accordance with the exemplary embodiment of the 25 invention are also ascribed identifiers and correspond to managed entities in the network management context. Naming conventions may be used in generating LSP tunnel identifiers 332. Selecting an LSP transport tunnel, the corresponding pair 340 of PE communications network nodes 102, and the LSP transport parameters 342 are displayed.

[66] Within each provisioned VPN service, and therefore in respect of each full mesh of LSP tunnels, actual traffic patterns may be inhomogeneous between the PE communications network node pairs; additional LSP tunnels being necessary between particular pair of PE network nodes to convey all the traffic, as shown in FIG. 2.

5 [67] The decommissioning 331 of an LSP transport tunnel may be made conditional on the non-existence of any VPN services actively commissioned therethrough. Updating 344 LSP transport tunnel parameters 342 may also be conditional on the non-existence of any VPN services actively commissioned through the corresponding LSP transport tunnel.

10 [68] A selected transport LSP full mesh may be decommissioned by interacting with button 346. Decommissioning an transport LSP full mesh may be conditional on the non-existence of any VPN services actively commissioned therethrough.

[69] FIG. 4 shows a state transition diagram corresponding to the full mesh transport LSP configuration and provisioning process described above.

15 [70] Therefore the full mesh transport LSP configuration means described, hides the physical layer complexities of the communications network 104 and provides a content transport provisioning framework for a multitude of L2 and L3 VPN services.

[71] In support of Pseudo-Wire (PW) and VPLS (L2VPN) service provisioning, signaling means must be provided between the group of PE network nodes 102 associated with a full transport LSP mesh to establish VPN connections therebetween. The signaling means includes targeted LDP sessions.

20 [72] In accordance with the exemplary embodiment of the invention, the functionality of the NMS 100 is further augmented with centralized full mesh targeted LDP signaling session establishment means. Provisioned transport LSP meshes and targeted LDP signaling session meshes are provisioned independently and are used together to enable PW and VPLS service provisioning.

25 [73] In accordance with an exemplary full mesh targeted LDP signaling session establishment method, operations management personnel selects a group of PE network nodes 102 to provision PW services therebetween for customer sites connected thereto,

targeted LDP signaling session commissioning commands are issued from the NMS 100 to each PE 102 in the group to effect the targeted LDP signaling session establishment therebetween.

[74] In selecting the group 612 of PE network nodes 102, operations management personnel, interacting with a Targeted LDP Signaling Session Mesh Configuration Module 600, having an exemplary interface schematically shown in FIG. 6, may be provided with a stylized map of the managed communications network 104 by interacting with a Show Network Map button 302 as described above.

[75] The NMS 100 may further query the NMS repository 120 for a list 304 of communications network nodes 102, by interacting with a Get Node List button 306 as described above.

[76] Alternatively, in selecting the group 312 of PE communications network nodes 102, having retrieved a list of communications network nodes 102, operations management personnel may be provided with filtering means 308 to identify PE communications network nodes 102 for inclusion 610 into the group 612 of PE network nodes 102.

[77] In order for the targeted LDP signaling sessions to be established between the PEs in the group, each PE communications network node 102 must also have LSR functionality. Having LSR functionality may be one of the selection criteria used in filtering as mentioned above (see “Is LSR” filter 308).

[78] Targeted LDP signaling sessions may be commissioned to convey signaling information therein in accordance with signaling parameters 614 including, but not limited to: Hello hold-time, Keep hold-time, Path Vector Limit (Pvl), etc. Signaling mesh parameters 614 may be specified and accepted by interacting with an Update Mesh Parameters button 616.

[79] By interacting with a Commission Signaling Mesh button 618, each LSR on a corresponding pair of PE network nodes 102 is instructed by the NMS 100 to establish targeted LDP signaling sessions therebetween relieving operations management personnel of doing so manually. $(N^*(N-1))/2$ bi-directional targeted LDP signaling

sessions are therefore commissioned in respect of each full signaling mesh in providing signaling information exchange between N PE communications network nodes 102.

[80] It is possible for a targeted LDP signaling session to already be commissioned between a particular pair of PE network nodes 102 prior to full targeted LDP signaling 5 session mesh commissioning, the targeted LDP signaling session already conveying signaling information. In accordance with an exemplary implementation of the exemplary embodiment of the invention, a process 800 of commissioning 618 a targeted LDP signaling session mesh shown in FIG. 8, may include determining 806 whether targeted LDP signaling sessions, themselves managed entities in the network 10 management context, already exist between PE communications network node pairs 802 and including 810 the existing targeted LDP signaling sessions in the signaling mesh being commissioned. An exemplary iterative (804, 812) process 800 is shown. Accordingly, the NMS full signaling mesh commissioning means 600, further includes functionality for discovering 806 actively commissioned targeted LDP signaling 15 sessions.

[81] In the following, targeted LDP signaling session establishment between PE network node pairs is assumed, although in practice, the path taken by signaling information conveyed in respect of a targeted LDP signaling session may vary in accordance with communications network conditions experienced thus providing 20 resilience therein.

[82] In accordance with the exemplary embodiment of the invention, the full mesh of established targeted LDP signaling session between a group of PE communications network nodes 612 represents a managed entity in the network management context. A signaling mesh identifier 622 is associated with the full signaling mesh managed entity 25 for further selection and reference, and a descriptive string 624 may be further associated with the managed signaling mesh entity. A status 626 may be ascribed to the managed signaling mesh entity and updated based on a multitude of related events. A Get Signaling Mesh List button 628 is provided to query the NMS database 120 for managed signaling mesh entities.

[83] Depending on the implementation and/or details of a particular deployment in support of specific VPN services, multiple targeted LDP signaling sessions terminating on interfaces on PE communications network nodes may be commissioned between particular PE network node pairs.

5 [84] Individual PE communications network nodes 102 may be added 610 to or removed 611 from a PE group 612, as well additional targeted LDP signaling sessions may be commissioned 630 between interfaces on PE network nodes 102 in the group 312, or decommissioned 631.

10 [85] With each PE network node 102 addition 610 to an existing group of N PE network nodes 612, N bi-directional targeted LDP signaling sessions are commissioned between N pairs of PE network nodes: between the newly added PE network node 102 and the existing N PE network nodes 102 in the group 612. With each PE network node 102 removal 611 from an existing group of N PE network nodes 612, N-1 bidirectional targeted LDP signaling sessions are decommissioned between N-1 pairs of PE network 15 nodes: between the PE network node to be removed and the remaining N-1 PE network nodes 102 in the group 612.

[86] The removal 611 of a PE network node 102 from the PE group may be conditional on the non-existence of any VPN services actively conveying signaling information therethrough. The restriction condition may be further limited to VPN 20 services terminating on the PE network node 102 to be removed 611, as the removed PE 102 may still convey signaling information in respect of managed VPN services as an intermediary MPLS communications network node 102.

[87] Targeted LDP signaling sessions also represent managed entities in the NMS management context and are tracked in the NMS repository 120. Each targeted LDP 25 signaling session has an associated identifier 632, a description 634, and an status 636. The NMS repository 120 may be queried, dependent on the active network management context, by interacting with a Get Signaling Session List button 638. If the get signaling session list button 638 is employed while a signaling mesh is selected, then only the targeted LDP signaling sessions of the selected signaling mesh are displayed, otherwise 30 all targeted LDP signaling sessions are displayed. Targeted LDP signaling sessions

commissioned in accordance with the exemplary embodiment of the invention are also ascribed identifiers and correspond to managed entities in the network management context. Naming conventions may be used in generating targeted LDP signaling session identifiers 632. Selecting a targeted LDP signaling session, the corresponding pair 640 5 of PE communications network nodes 102, corresponding pair of interfaces 641, and the targeted LDP signaling session parameters 642 are displayed.

[88] The decommissioning 631 of a targeted LDP signaling session may be made conditional on the non-existence of any VPN services actively conveying signaling information therethrough. Updating 644 targeted LDP signaling session parameters 642 10 may also be conditional on the non-existence of any VPN services actively conveying signaling information through the corresponding targeted LDP signaling session.

[89] A selected full mesh of targeted LDP signaling sessions may be decommissioned by interacting with button 646. Decommissioning a full mesh of targeted LDP signaling sessions may be conditional on the non-existence of any VPN services actively 15 conveying signaling information therethrough.

[90] FIG. 7 shows a state transition diagram corresponding to the full targeted LDP signaling session mesh configuration and provisioning process described above.

[91] Therefore the full targeted LDP signaling session mesh configuration means described, hides the physical layer complexities of the communications network 104 and 20 provides a signaling information transport provisioning framework for a multitude of Layer-2 and Layer-3 VPN services.

[92] In accordance with the exemplary embodiment of the invention, the full mesh LSP content transport provisioning and full mesh targeted LDP signaling session provisioning effected in a centralized network management context enables service 25 providers to: assure network resiliency, assure service quality, and provide accounting in respect of Layer-2 and Layer-3 VPN services irrespective of client-side deployed infrastructure while leveraging installed infrastructure in the service provider communications network core.

[93] Further Layer-2 and Layer-3 VPN service provisioning details are provided herein below in respect of particular infrastructure employed.

[94] In accordance with a VPWS services provisioning, Ethernet network interworking over MPLS is described.

5 [95] In accordance with the exemplary embodiment of the invention, Layer-2 Ethernet packets are enabled to traverse an MPLS service provider communications network 104 connecting two serviced customer Ethernet LAN segments 20 together. Packets received on an exemplary ingress 1Gb Ethernet (GigE) interface of a PE network node 102 and are transmitted out of the originating PE network node 102 via an
10 MPLS capable interface, such as, but not limited to Packet over SONET (PoS), GigE, or ATM. Making reference to FIG. 9 two-deep MPLS label stack is employed as defined in IETF draft-martini-ethernet-encap-mpls-01.txt (draft-martini), which is incorporated herein by reference, where the inner label corresponds to the Layer-2 Ethernet connection and the outer label is used to forward MPLS encapsulated packets through
15 the service provider communications network 104 onto the appropriate transport LSP/mesh. Label Distribution Protocol Downstream-Unsolicited (LDP-DU) signaled LSPs (S-LSPs) may be employed in provisioning the inner (label) LSP. RVSP-TE signaling is used to provision the outer (label) LSP. Labeled packets travel through the MPLS service provider communications network 104 to an MPLS capable interface
20 (PoS, GigE, or ATM) to the egress PE network node 102. The packets then exit the egress PE network node 102 via a GigE interface and are conveyed to the destination customer Ethernet LAN segment 20.

[96] As described in draft-martini, a GigE port or VLAN interface is configured for network interworking, remote router, and Martini parameters. This initiates the
25 signaling of the inner label through an LDP advertisement of the reachability of an Ethernet interface through an allocated MPLS label. When a match, of configured, Martini parameters is found between a configured interface and an advertised interface, the Layer-2 VPN connection may be configured and content traffic may be conveyed. All packets received on the ingress GigE interface are transported to the remote egress
30 GigE interface.

[97] In order for Ethernet VPWS connections to be made, a signaled MPLS tunnel must exist between the PE communications network nodes 102 that attach to the customer equipment 22. To establish a VPWS connection (pseudo-wire) over a signaled tunnel LSP, specific pseudo-wire configurations must be performed on both PE
5 LSP routers 102.

[98] When VPWS connections terminate at CE's 20 attached by dissimilar layer 2 technologies (e.g. ATM and GigE), the service offering requires conversion of packets from one Layer-2 technology to another, known as service interworking. This packet conversion or service interworking is performed at the PE network nodes 102 in the
10 service provider's communications network 104.

[99] As technology advances and the cost of higher bandwidth Ethernet equipment become more economical compared to OC-3 and OC-12 ATM equipment. Existing point-to-point networking using ATM point-to-point services may be migrated to use 100baseT and GigE interfaces at access point PE network nodes 12. As such, service
15 interworking may offer a way to upgrade existing ATM point-to-point network support using VPWS, as shown in FIG. 10, creating a point-to-point pseudo-wires between ATM and GigE interfaces. Layer-2 Ethernet packets are enabled to traverse an MPLS service provider communications network 104 and connect to a Layer-2 PE switch 102 or CE router 22 supporting ATM. In order for Ethernet/ATM service interworking
20 VPWS connections to be made, a signaled MPLS tunnel must exist. Also, the ATM Virtual Path Identifier / Virtual Connection Identifier (VPI/VCI) must be encapsulated using ATM Adaptation Layer-5 (AAL5).

[100] Although, in the above extensive reference was made to PW services and L2VPN services in general, the invention is not limited thereto, implementations of the
25 exemplary embodiment of the invention may also include provisioning support for Virtual Private Wire Services (VPWS) and Virtual Private LAN Service (VPLS), as well as Layer-3 VPN services.

[101] Further the invention is not limited to any transport LSP type. A variety of transport LSPs may be employed including, but not limited to: Label Distribution

Protocol LSPs, Resource Reservation Protocol-Traffic Engineering (RSVP-TE) LSPs,
Labeled LSP, etc.

[102] The embodiments presented are exemplary only and persons skilled in the art
would appreciate that variations to the above described embodiments may be made
5 without departing from the spirit of the invention. The scope of the invention is solely
defined by the appended claims.